Journal of Economic Behavior & Organization 132 (2016) 56–65

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

How should economists model climate? Tipping points and nonlinear dynamics of carbon dioxide concentrations☆

Jean-Paul Chavas, Corbett Grainger*, Nicholas Hudson

University of Wisconsin, Madison, United States

ARTICLE INFO

Article history:
Received 14 September 2015
Received in revised form 27 January 2016
Accepted 29 January 2016
Available online 19 February 2016

Keywords:
Tipping points
Threshold quantile autoregression
Nonlinear dynamics
Climate change

ABSTRACT

Economists modeling climate policy face an array of choices when modeling climate change, including the role of uncertainty/ambiguity, irreversibility, and tipping points. After filtering out estimated cycles due to orbital climate forcing, we use a threshold quantile autoregressive model to characterize anomalies in atmospheric CO_2 concentrations. We then test for local instability and tipping points, and we characterize the stationary distribution of anomalies. We find evidence of nonlinear dynamics, tipping points and a non-normal stationary distribution.

1. Introduction

The economic analysis of climate change is complicated by the vast array of policy choices and the complexity of the natural system and the associated interactions with the economy (Nordhaus, 2008). While models have made great strides forward over the past several decades, economists still face difficult choices of how to characterize the natural system. For example, whether climate change is slow or abrupt, whether the distribution of damages has "thick" or "thin" tails (Weitzman, 2009), the extent to which climate change is irreversible (Barrett and Dannenberg, 2012), and the existence of tipping points (Lemoine and Traeger, 2014) all affect the optimal policy response. In this paper we examine the paleoclimate data going back roughly 400,000 years to characterize historical climate dynamics, look for evidence of (ir)reversibility, and present evidence of tipping points and critical slowing. We use a novel econometric approach—a Threshold Quantile Autoregressive model—that is ideal for characterizing this type of dynamic system. The results of this analysis informs the choices made in constructing economic models of climate change.

Determining the existence of tipping points in any dynamic system is an empirical challenge for economists and natural scientists alike. Examples abound: ecological systems (Carpenter et al., 1985), coupled human and natural systems (Liu et al., 2007), or economic phenomena such as business cycles or unemployment (Galvao et al., 2011). Our application utilizes paleoclimate ice core data to test for evidence of local instability in carbon dioxide (CO_2) concentrations. Empirically identifying the presence of tipping points is a nontrivial exercise. A tipping point may exist in a zone of instability surrounded by zones of stability: this implies an "escape" from the unstable zone around the tipping point and a movement toward the

☆ We thank participants of the "Thresholds, Tipping Points and Random Events in Dynamic Economic Systems" workshop at The Howard H. Baker Jr. Center for Public Policy. We are especially thankful to Jacob LaRiviere, Derek Lemoine, Christian Traeger, Chuck Mason, Charles Sims, David Finnoff and Bill Neilson.

* Corresponding author. Tel.: +1 608 262 3651.
E-mail address: corbett.grainger@wisc.edu (C. Grainger).

http://dx.doi.org/10.1016/j.jebo.2016.01.013
0167-2681/

surrounding stable zones. Thus, our estimates of local instability can be interpreted as evidence of tipping points in CO_2 dynamics.

A relevant issue is the potentially irreversible nature of temperature increases following shocks to atmospheric greenhouse gas concentrations, raising concerns about possibly catastrophic implications of climate change. Recent research using uncovered geological evidence (in sediment records) demonstrates a precedent of doubling atmospheric CO_2 concentrations 55 million years ago which led to rapid oceanic acidification and global temperature increases of 5 °C over a relatively short geological time period (Wright and Schaller, 2013). If climate change is irreversible, steps toward mitigation and adaptation should be taken early if altering climate is costly. Alternatively, if it is not irreversible, the optimal policy may not require aggressive early action. Our empirical results help shed light on the nature and speed of dynamic adjustments in CO_2, including how historical CO_2 concentrations have evolved after local maxima or minima.

The existence of tipping points or evidence of reversibility are both critical to our understanding of climate change because the optimal policy response may vary depending on whether there are tipping points, the location(s) of tipping points, and whether change is reversible. Once a threshold has been crossed, alternative policy responses such as geoengineering have been proposed as solutions. In this context, for example, Barrett and Dannenberg (2012) argue that the existence of tipping points can turn the international environmental agreement game from a prisoner's dilemma into a coordination game.

In addition, the shape of the probability distribution of historical CO_2 concentrations is of interest as it can affect the optimal policy. While uncertainty and learning have been incorporated into climate models for decades (Kelly and Kolstad, 1999), a recent debate (e.g., Weitzman, 2009) suggests that the shape of the distribution is critical. Specifically, the presence of "thick tails" would suggest that the likelihood of rare (and possibly catastrophic) events would occur more frequently than previously thought. We aim to help inform this debate by characterizing the stationary distribution of CO_2 based on data from the past 400,000 years.

We use a novel econometric approach to examine the dynamics of historical climate. The model, a Threshold Quantile Autoregressive (TQAR) model (see e.g., Galvao et al., 2011; Chavas, 2015) represents the distribution of CO_2 conditional on its previous history. We allow the lag effects to vary across quantiles of the distribution as well as with previous levels. Importantly, this allows us to capture nonlinear dynamics and to characterize how lag effects of atmospheric CO_2 concentrations influence current levels at different points in the distribution. This, in turn, gives us a basis to investigate the presence of thresholds and tipping points.

In what follows, we examine the nature of CO_2 dynamics applied to the Vostok ice core data series, which contains CO_2 concentrations for the past 400,000 years, observed at roughly 1000 year intervals. We focus on CO_2 rather than reconstructed temperature from climatic models because the historical atmospheric constituents are measured directly.[1] We demonstrate that our approach can characterize some well-known long-run cycles in climate known as Milankovitch cycles (Jansen et al., 2007). Loosely explained, these cycles, which are well understood by climate scientists, are caused by orbital irregularities (eccentricity), fluctuations in the axial tilt (obliquity), and axial directionality (precession).[2] We then examine the dynamics of anomalies (defined as deviations from Milankovitch cycles) using a TQAR model. The estimated model exhibits slow-moving nonlinear dynamics.

Our empirical results indicate local instability in CO_2 concentrations, indicating the presence of tipping points. We then simulate the dynamic model using Markov Chains to obtain the stationary distribution and find that the stationary distribution is non-normal. These two findings indicate that economists modeling climate change policy should account for the presence of tipping points and reversibility.

2. Preliminary analysis of climate cycles

Paleoclimatologists have studied the cyclical nature of climate in several settings, including geological deposits, atmospheric constituents trapped in ice cores, and using fossil records to infer changes in historical temperatures and CO_2 concentrations (Solomon, 2007). Scientific consensus typically attributes the so-called Milankovitch Cycles to be a major determinant in long-run climate variation (Milanković, 1941). In the early 20th century, Milankovitch hypothesized that climate was determined by changes in the Earth's orbit. An excellent summary of empirical findings is in Kaper and Engler (2013). In addition to three well-known "Milankovitch Cycles", a fourth cycle has been identified empirically (Petit et al., 1999).[3] After the cycles are estimated, we de-cycle the data to identify anomalies.

Following Chavas et al. (2015), in order to capture the periodicity of the CO_2 evolution over time we estimate the following equation:

$$CO_2(t) = a + \sum_{i=1}^{N} b_i \cdot \sin\left(k_i + \frac{2\pi}{c_i} t\right) + e_t, \tag{1}$$

[1] The deuterium, δD_{ice}, and isotopic oxygen, $\delta^{18}O$, levels in sea ice are typically used to calculate local temperature changes from ice core data in a procedure outlined by Lorius and Merlivat (1975).

[2] In addition to the three main effects documented in the literature, we also identify a fourth cycle with a shorter period. This fourth cycle is not common in the literature, though it has been detected using other statistical methods such as spectral analysis (e.g., Petit et al., 1999).

[3] Also see Chavas et al. (2015).

Table 1
Milankovitch cycle estimates.

Wavelength bounds in kyrs	Cycle estimate in kyrs	Cycles in kyrs as computed by Petit et al. (1999)
60–120	98.200*** (1.000)	100
0–60	39.725*** (0.161)	43
0–60	28.716*** (0.086)	24
0–60	17.890*** (0.033)	19

Notes: Cycles were determined using the R package "genalg" which randomly selected values for each of the parameters from a given parameter space. See Chavas et al. (2015) for details. Standard errors are calculated using a variance–covariance matrix for nonlinear LS estimators as in Judge et al. (1988) pp. 640–43.
*** Asterisks denote significance at the 0.1 percent level.

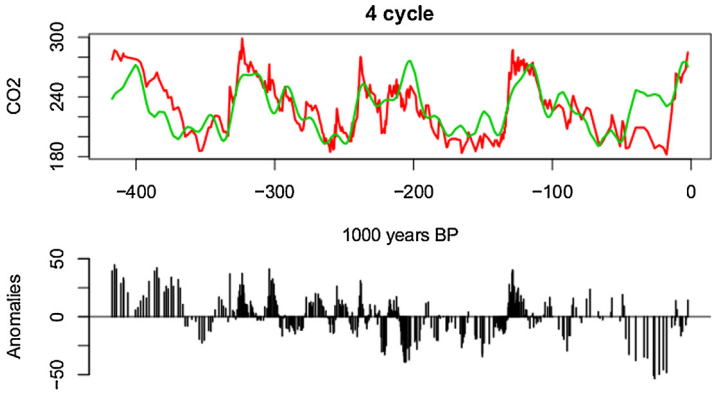

Fig. 1. Actual vs. predicted Vostok CO_2 concentrations. Notes: CO_2 concentrations in p.p.m.v. as collected from 417,160 to 2342 years before present (BP). Actual values are in red. Predicted values (in green) are estimated as described in the text. The bottom panel shows the anomalies (residuals). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

where $CO_2(t)$ is the CO_2 concentration from the Vostok ice core data, t is time (1000 years), b_i is the amplitude, k_i is the horizontal shift, c_i is the period of the ith cycle, $i = 1, \ldots, N$, N is the number of cycles, and e_t is an error term distributed with mean zero and finite variance.

The estimation of the cycles in Eq. (1) can prove difficult. Indeed, using the least squares method, the estimation of the parameters in (1) involves many local minima. To deal with this issue, Eq. (1) was estimated in two steps. First, we use a genetic algorithm (Lucasius and Kateman, 1993, 1994) to identify parameter estimates yielding global least-squares minima. Second, hypothesis testing was conducted to evaluate the relevant number cycles N in Eq. (1). The analysis presents strong evidence indicating the presence of four cycles. See Chavas et al. (2015) for details. The econometric estimates of Eq. (1) are presented in Table 1.

The parameter estimates reported in Table 1 indicate that the four cycles have periods of 98,200, 39,725, 28,716 and 17,890 years, respectively. These cycle periods have been investigated extensively in previous research (e.g., Petit et al., 1999). For example, the first cycle (with a period around 100,000 years) is well-identified and has been associated with the eccentricity of earth's movements, and the second cycle (with a period of about 40,000 years) has been associated with obliquity of the earth's axial tilt. Our estimates of CO_2 cycles do not control for other physical forcings, so there are likely omitted variables; we proceed with that caveat in mind. Fig. 1 shows the fit of the preferred model along with the data.

Going beyond the Milankovitch cycles, we are interested in the dynamic process affecting the evolution of CO_2 over time. Our analysis proceeds evaluating the dynamics of the error term e_t in Eq. (1). Consistent with the literature on climate change, the error term e_t is defined as an "anomaly" measuring the deviation of actual CO_2 from its expected value along our estimated cycles. A preliminary analysis of anomalies is presented in Table 2. It starts with a standard autoregressive model of order m, AR(m), expressing current anomaly e_t as a linear function of past anomalies ($e_{t-1}, \ldots, e_{t-m}$), where we invoke the "finite m memory" assumption. Since the data involve unequally spaced observations, we introduce the variables $d_{t-i} = [(year_{t-i+1} - year_{t-i})/1000 - 1]$, which reflects the of time difference between the observation $t - i + 1$ and $t - i$, $i = 1, \ldots, m$. The variables d capture the effects of unequally spaced observations on CO_2 dynamics.

Table 2 reports the estimates of two AR(m) models: AR(1) and AR(2). We also consider the possibility that the Autoregressive coefficient in the AR(m) models may not be constant. We partition the set of anomalies into three subsets defining three regimes: the regime S_1 when the anomaly e is in the interval $(-\infty, s_1)$, the regime S_2 when the anomaly is in the interval $[s_1, s_2)$, and the regime S_3 when the anomaly is in the interval $[s_2, \infty)$, where s_1 and s_2 are threshold points corresponding

Table 2
Preliminary analysis of anomalies.

Parameters	AR(1)	TAR(1)	AR(2)	TAR(2)
Intercept	−0.083	0.832	−0.079	0.543
	(0.461)	(0.581)	(0.461)	(0.604)
e_{t-1}	0.839***	0.950***	0.811***	0.954***
	(0.032)	(0.106)	(0.055)	(0.139)
e_{t-2}			0.024	0.011
			(0.055)	(0.138)
d_{t-1}	−0.003	−0.005	−0.012	−0.015
	(0.021)	(0.022)	(0.022)	(0.023)
d_{t-2}			0.023	0.027
			(0.022)	(0.023)
$(e_{t-1}-s_1)^* r_{1,t-1}$		−0.038		0.037
		(0.161)		(0.198)
$(e_{t-1}-s_2)^* r_{3,t-1}$		−0.323		−0.388*
		(0.167)		(0.197)
$(e_{t-2}-s_1)^* r_{1,t-2}$				−0.142
				(0.195)
$(e_{t-2}-s_2)^* r_{3,t-2}$				0.077
				(0.195)
Standard error	8.474	8.668	8.752	8.673
R square	0.708	0.715	0.710	0.718
Adjusted R square	0.707	0.712	0.706	0.712
BIC	2610.856	2613.998	2621.000	2633.941

Note: Standard errors are presented in parentheses below the parameter estimates. Standard errors are bootstrapped using the paired bootstrap method. Asterisks indicate the significance level: ***, ** and * represent significance at the 0.1%, 1%, and 5% levels, respectively.

respectively to the 1/3 and 2/3 quantile of the distribution of anomalies.[4] In this context, we define the dummy variables $r_{j,t-i} = \mathbf{1}\left\{e_{t-i} \in S_j\right\}$ where $r_{j,t-i}$ is equal to 1 when e_{t-i} is in the jth regime. Then, introducing the variables $[(e_{t-i}-s_1)^* r_{1,t-i}]$ and $[(e_{t-i}-s_2)^* r_{3,t-i}]$ into the AR(m) model generates a Threshold Autoregressive model, TAR(m). A TAR(m) model allows the autoregression parameters to vary across regimes. The estimates of the TAR(m) models, TAR(1) and TAR(2), are also reported in Table 2.

The models reported in Table 2 show statistical evidence of dynamics in the anomalies. Indeed, the coefficient of e_{t-1} is statistically significant in all four models, providing evidence that temporal adjustments to shocks are slow. That is, the coefficients indicate that lagged anomalies are a significant determinant of current shocks. Interestingly, this coefficient varies between 0.811 in the AR(2) model to 0.954 in the TAR(2) model. In addition, the dominant root of the four dynamic systems is always less than 1, indicating that all estimated models in Table 2 exhibit stable dynamics. We will obtain additional information on this issue from our TQAR model (see below).

Which of the four models in Table 2 provides a better representation of the dynamics of anomalies? Conducting hypothesis testing can help this evaluation. Using an F-test of the joint significance of e_{t-2} and d_{t-2} in the AR(2) model (i.e., comparing AR(1) and AR(2)), we obtain a F-value of 0.808, with a corresponding p-value of 0.446. Thus, we fail to reject the AR(1) model (compared to AR(2)). Comparing AR(1) and TAR(1), we obtain a F-value of 4.310, with a p-value of 0.014, providing a strong rejection of AR(1) in favor of TAR(1). This provides statistical evidence that the autoregression parameters are not constant and vary across regimes. Finally, comparing TAR(1) and TAR(2), we obtain a F value of 0.885, with a p-value of 0.473, indicating that we fail to reject the TAR(1) model. The adjusted R^2 and BIC provide weak support for the single lag models over the two-lag models. In order to test for the presence of thresholds and reversibility, we require a TQAR model, which we address in the following section.

3. TQAR analysis

We now consider a Threshold Quantile Autoregression (TQAR) model applied to CO_2 anomalies. A discussion of TQAR models is presented in Galvao et al. (2011) and Chavas (2015). Having found support for a TAR(1) specification, we examine below a TQAR(1) specification. Let $F(u|e_{t-1}) = Prob[e_t \le u|e_{t-1}]$ be the distribution function of e_t conditional on e_{t-1}. The conditional quantile function is defined as the inverse function $Q(q|e_{t-1}) \equiv \inf_c \left\{c|F(c|e_{t-1}) \ge q\right\}$ for $q \in (0, 1)$. This includes the conditional median when $q = 0.5$. The distribution function $F(\cdot)$ and the quantile function $Q(\cdot)$ provide general characterizations of the dynamics of the anomalies e_t.

[4] The analysis presented below was also conducted using different threshold points (s_1, s_2). Our key qualitative findings (including local instability) were found to be robust to the choice of (s_1, s_2).

Table 3
Estimation of TQAR(1) model for selected quantiles.

Quantile	$q=0.1$	$q=0.3$	$q=0.5$	$q=0.7$	$q=0.9$
Intercept	−7.782***	−3.050***	0.128	3.989***	10.210***
	(0.895)	(0.520)	(0.627)	(0.726)	(1.318)
e_{t-1}	0.902***	0.940***	0.904***	0.931***	1.307***
	(0.182)	(0.091)	(0.115)	(0.135)	(0.258)
d_{t-1}	0.057	0.021	−0.019	0.012	0.062
	(0.050)	(0.024)	(0.044)	(0.055)	(0.079)
$(e_{t-1}-s_1)*r_{1,t-1}$	0.036	0.085	0.071	−0.038	−0.611
	(0.253)	(0.142)	(0.183)	(0.238)	(0.454)
$(e_{t-1}-s_2)*r_{3,t-1}$	−0.394	−0.389*	−0.111	−0.231	−0.708
	(0.294)	(0.174)	(0.227)	(0.201)	(0.371)

Note: Standard errors are presented in parentheses below the parameter estimates. Standard errors are bootstapped using the paired bootstrap method. Asterisks indicate the significance level: ***, ** and * represent significance at the 0.1%, 1%, and 5% levels, respectively.

Table 4
Hypothesis testing under the TQAR(1) model.

Hypothesis test	F-value	Deg of freedom	p-value
Testing TQAR(1) vs. QAR(1) for selected quantiles	1.772**	(32, 3217)	0.005
$q=0.1$	1.361	(2, 356)	0.258
$q=0.2$	6.999***	(2, 356)	0.001
$q=0.3$	8.587***	(2, 356)	0.000
$q=0.4$	3.364*	(2, 356)	0.036
$q=0.5$	0.728	(2, 356)	0.484
$q=0.6$	3.038*	(2, 356)	0.049
$q=0.7$	1.658	(2, 356)	0.192
$q=0.8$	1.766	(2, 356)	0.173
$q=0.9$	1.641	(2, 356)	0.195

Note: Standard errors are bootstrapped using the paired bootstrap method. Asterisks indicate the significance level: ***, ** and * represent significance at the 0.1%, 1%, and 5% levels, respectively.

When evaluated in the qth quantile, we consider the TQAR(1) model specified as

$$Q(q|e_{t-1}) = \alpha_q + e_{t-1} * \beta_q + d_{t-1} * \gamma_q + (e_{t-1} - s_1) * r_{1,t-1} * \beta_{1q} + (e_{t-1} - s_2) * r_{3,t-1} * \beta_{3q}, \quad (2)$$

where the variables (d, s, r) are defined as in the previous section. As in all quantile regression models (e.g., Koenker, 2005), the parameters $(\alpha_q, \gamma_q, \beta_q, \beta_{1q}, \beta_{2q})$ are allowed to vary across quantiles q. This is more general than the standard autoregressive (AR) models. In addition, the specification (2) allows for three dynamic regimes. The parameter β_q in Eq. (2) captures the marginal effect of e_{t-1} when e_t is located in the second quantile (regime S_2) of its distribution. But this marginal effect becomes $(\beta_q + \beta_{1q})$ when e_t is in its first quantile (regime S_1) and $(\beta_q + \beta_{3q})$ when e_t is in its third quantile (regime S_3).

The TQAR model (2) includes a number of well-known models as special cases. First, (2) reduces to the Quantile Autoregression (QAR) model proposed by Koenker and Xiao (2006) when $\beta_{1q} = \beta_{3q} = 0$, i.e., when the parameters are the same in all three regimes. Note this still allows the parameters to vary across quantiles q. If in addition, the parameters are the same across all quantiles q, then (2) reduces to a standard Autoregression (AR) model where the regression parameters are treated as constant.

The TQAR model (2) is flexible as it allows for two forms of nonlinear dynamics. First, (2) allows for the effects of e_{t-1} to vary with quantile q. This is of interest if dynamic adjustments depend on the nature of current shocks. Second, (2) allows for the effects of e_{t-1} to vary depending on the regime where e_{t-1} is located. This is of interest if dynamic adjustments depend on the nature of past shocks. Neither of these properties is present in the standard autoregression (AR) model; and only the first characteristic is present in the quantile autoregressive (QAR) model. Being linear in the parameters, model (2) can be estimated as a quantile regression model. The estimation method and the statistical properties of its estimator have been analyzed in previous literature (e.g., Koenker, 2005; Koenker and Xiao, 2006).

The TQAR model allows one to characterize nonlinear dynamics because it allows adjustments to vary with the nature of past and current shocks. Furthermore, in contrast to the TAR model, the TQAR model allows the effects vary by quantile. The presence of tipping points can be characterized by zones of instability surrounded by zones of stability, which the TQAR model can identify.

The estimation of a TQAR(1) model (2) is reported in Table 3 for selected quantiles, $q = (0.1, 0.3, 0.5, 0.7, 0.9)$. In Table 3, the coefficient of the lagged anomaly e_{t-1} is found to be statistically significant across all reported quantiles. Again, this provides strong evidence of dynamics in anomalies. In addition, Table 4 shows that, at least for the quantile $q = 0.3$, the Autoregression parameters are not constant across regimes.

How appropriate is the TQAR(1) model specification? To help answer this question, hypothesis testing is conducted on the TQAR(1) model. First, we test the joint significance of whether the parameters across quantiles $q = (0.1, 0.2, \ldots, 0.8, 0.9)$. Testing the joint significance of these parameters, the F-statistic is 1.77, which is significant at the 1 percent level. This

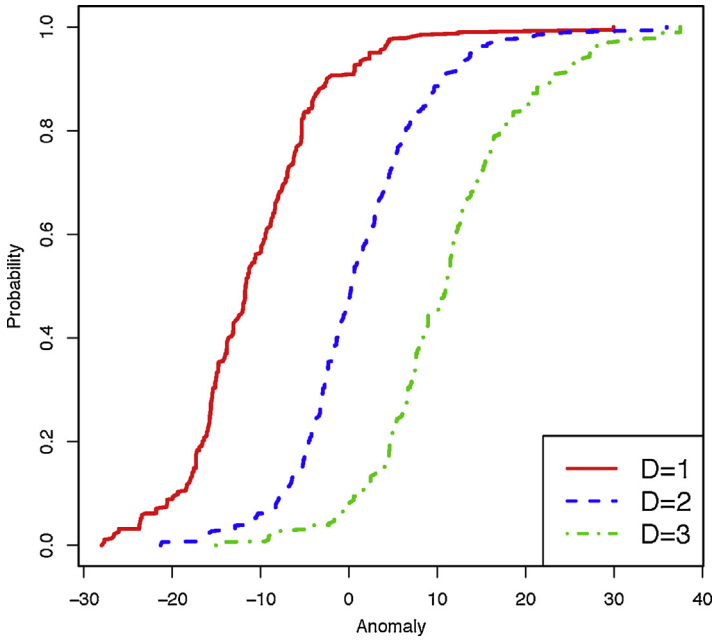

Fig. 2. Distribution function of anomalies.

provides evidence that dynamics vary across quantiles, thus supporting our quantile approach. Second, for selected quantiles, we test the null hypothesis that the threshold parameters $(e_{t-1} - s_1)^* r_{1,t-1}$ and $(e_{t-1} - s_2)^* r_{3,t-1}$ are jointly significant; that is, we test whether parameters are constant across regimes. This amounts to testing the TQAR(1) model versus a QAR(1) model. These test results are reported in Table 4. They show strong evidence against QAR(1) and in favor of the TQAR(1) when $q = (0.2, 0.3, 0.4, 0.6)$. These results support the argument that the TQAR(1) model provides an appropriate and flexible representation of the dynamics of CO_2 anomalies. They document the presence of significant nonlinear dynamics in CO_2 anomalies, with dynamics varying with both current shocks and past shocks. This identifies the presence of asymmetric dynamic adjustments in CO_2 evolution across regimes.

Next, we investigate the implications of our TQAR(1) estimates. The rest of our analysis involves estimates of the TQAR(1) model after the model has been re-estimated for all quantiles $q \in (0, 1)$. First, using the parameter estimates for all quantiles, we simulate the predicted distribution function of the anomalies under three scenarios: $D = 1, 2, 3$. The scenarios depend on the value taken by e_{t-1} and its location on the distribution function of anomalies: scenario 1 ($D = 1$) is when e_{t-1} is in regime S_1 (i.e., in the lower tertile); scenario 2 ($D = 2$) is when e_{t-1} is in regime S_2 (i.e., in the middle tercile); and scenario 3 ($D = 3$) is when e_{t-1} is in regime S_3 (i.e., in the upper tercile). The estimated distribution function of the anomalies is presented in Fig. 2 under these three scenarios. As expected, in the presence of slow adjustments, the distribution function of CO_2 anomalies shifts to the right as one moves from scenario 1 ($D = 1$) to scenario 2 ($D = 2$) to scenario 3 ($D = 3$). All three estimates show that the simulated distributions of anomalies exhibit a range from about −30 to +30.

As noted above, in the TQAR(1) model, the estimated coefficient of the lagged anomaly e_{t-1} varies both across quantiles q and across regimes (S_1, S_2, S_3). Recall that regime S_1 occurs when the anomaly e_{t-1} is in the lower quantile [0, 1/3) of its distribution, regime S_2 when the anomaly is in the quantile [1/3, 2/3), and regime S_3 when the anomaly is in the upper quantile [2/3, 1]. Also note that, in the TQAR(1) model, the marginal effect of e_{t-1}, $\partial e_t/\partial e_{t-1}$, provides useful information about the nature and speed of dynamic adjustments. In this context, the logarithm of $|\partial e_t/\partial e_{t-1}|$ measures the rate of divergence of nearby dynamic trajectories of anomalies. It follows that the dynamics are locally stable (unstable) if $|\partial e_t/\partial e_{t-1}|$ is less than 1 (greater than 1). Noting that $|\partial e_t/\partial e_{t-1}|$ is the root of the dynamic system, local stability (instability) of the system occurs when its root is less than 1 (greater than 1). Under nonlinear dynamics, this evaluation is local and depends on the evaluation point. Applied to Table 3, we obtain important results. In Table 3, when e_{t-1} is in regime S_2, the estimated coefficient of e_{t-1} in the $q = 0.9$ quantile is 1.307, indicating a zone of local instability. In this zone, nearby trajectories diverge as anomalies tend to escape from this zone. Table 3 also shows that the coefficients of e_{t-1} are typically less than one in other quantiles, documenting zones of stability where nearby trajectories of anomalies converge.

We evaluated the root of the dynamic system (as measured by $|\partial e_t/\partial e_{t-1}|$ for all quantiles and all regimes. The results are presented in Fig. 3 which shows some interesting patterns. First, the root tends to be less than 1 and smaller in regime S_3, indicating that dynamic adjustments are locally stable when e_{t-1} is in the upper quantile of its distribution. Second, when e_{t-1} is in regime S_1, the root is close to 1 for all quantiles $q \in (0, 0.6)$, but it becomes much smaller than 1 when q is around 0.9. Thus, in regime S_1, the dynamics appears to be locally stable but only in the upper tail of the distribution. Third, when e_{t-1} is in regime S_2, the root is close to 1 for all quantiles $q \in (0, 0.8)$, but it becomes greater than 1 (corresponding to local

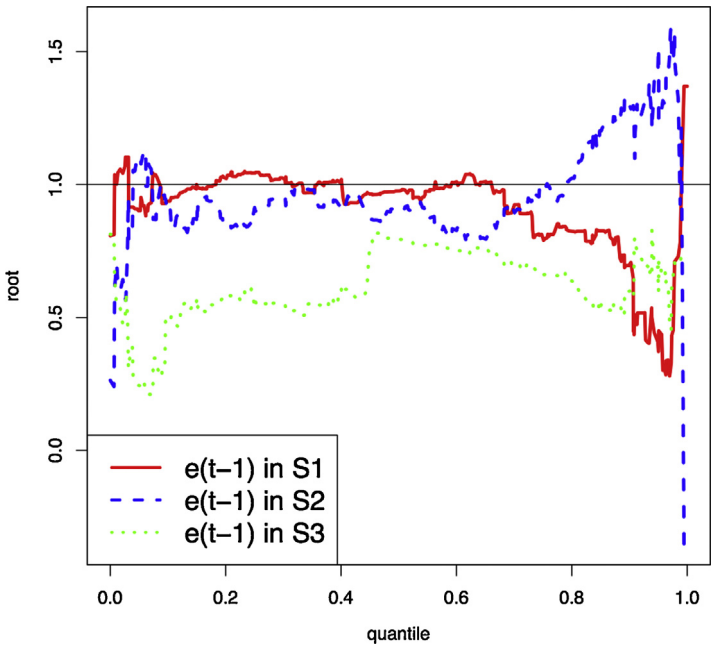

Fig. 3. Root for TQAR(1).

instability) when q is around 0.9. This reflects that dynamic adjustments tend to be qualitatively different across regimes, with evidence of dynamic local instability especially under regime S_2 and when q is around 0.9. These results reflect the presence of nonlinear dynamics, as the dynamics of anomalies vary with both current shocks and past shocks.

Importantly, our estimated TQAR model finds evidence of a zone of instability surrounded by zones of stability. This identifies the existence of tipping points. Indeed, a tipping point arises when dynamics moves the system away from a particular zone. This occurs when a zone of instability is surrounded by zones of stability, implying an "escape" from the unstable zone around the tipping point and a movement toward the surrounding stable zones. In this context, our estimates of local instability can be interpreted as evidence of tipping points in CO_2 dynamics. Our analysis indicates the presence of a tipping point when the previous shock e_{t-1} is in regime S_2 followed by positive shocks for e_t (i.e., when e_t is located around quantile 0.9).

4. Simulations

The TQAR model estimated in the previous section provides a representation of nonlinear dynamics for CO_2 anomalies. What does our analysis imply for the evolution of CO_2? To gain additional insights, we proceed by simulating our estimated model. This is done through discretizing the state space and rewriting our TQAR(1) model as a Markov chain model, which is then used to simulate the long-run evolution of the distribution of anomalies. The discretization is done by decomposing the observed range of CO_2 anomalies into 80 equally-spaced intervals, with each interval corresponding to one point in the Markov chain representation. The Markov chain model takes the form

$$P_t = AP_{t-1} \tag{3}$$

where $P_t = \begin{bmatrix} P_{1t} \\ P_{2t} \\ \vdots \end{bmatrix}$ is a (80×1) vector of probabilities where P_{jt} is the probability of being at point j at time t, and where $A = \begin{bmatrix} a_{11} & a_{12} & \cdots \\ a_{21} & a_{22} & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}$ is a (80×80) transition probability matrix where a_{ij} is the probability of moving from point j at time $t-1$ to point i at time t such that $\sum_i a_{ij} = 1$ for all j.

We evaluate the Eigenvalues of the matrix A. The 5 largest Eigenvalues are real and take the values (1, 0.863, 0.755, 0.643, 0.531). Finding that the largest Eigenvalue (1) is unique indicates that the Markov chain has a unique stationary distribution (Meyn and Tweedie, 2012). Furthermore, finding that the second largest Eigenvalue is 0.863 demonstrates that dynamic adjustments are slow.

Table 5
Transition probabilities.

Transition probabilities p_{ij}		Initial regime		
		S_1	S_2	S_3
Transition regime	S_1	0.776	0.159	0.019
	S_2	0.209	0.622	0.227
	S_3	0.015	0.219	0.754

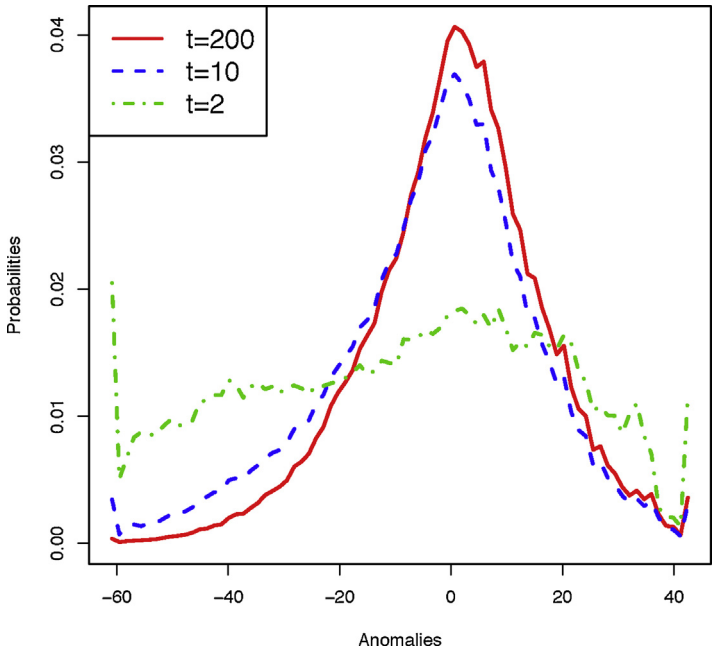

Fig. 4. Simulated probabilities.

Next, we simulate the evolution of the probability function using a uniform distribution as starting values, knowing the simulated probabilities will eventually converge to a stationary probability function. Evaluated at stationary probabilities, we calculated the probabilities of transitioning between the three regimes (S_1, S_2, S_3). This gives a 3×3 transition matrix, where an element p_{ij} is the probability of moving from regime S_i to regime S_j, such that $\Sigma_j p_{ij} = 1$ for $i = 1, 2, 3$. This transition matrix is reported in Table 5. Table 5 demonstrates a high probability of staying within any given regime. But it also shows that there is a positive probability of switching from any regime j to regime i.

Finally, the speed of convergence of the probabilities of the anomalies to their stationary values is presented in Fig. 4 after 2000 years ($t = 2$), 20,000 years ($t = 10$) and 200,000 years ($t = 200$). Fig. 4 illustrates that the convergence to the stationary distribution is slow. For example, only a fraction of the eventual probability adjustments take place during the first 2000 years. Second, Fig. 4 shows that the long-term probability function of anomalies has a unique maximum around 0, but it is skewed to the left. Testing whether the probability function is Normally distributed, we used the Shapiro–Wilk Normality Test (Shaphiro and Wilk, 1965). The test statistic is $W = 0.848$, with a p-value less than 0.001, showing strong statistical evidence against the Normal distribution. Thus, the estimated non-linear dynamics reported above contributes to a departure from Normality and to the presence of left-skewness. The implications of our findings are further explored next.

5. Discussion

We evaluated the dynamics of anomalies defined as deviations of observed CO_2 concentrations from estimated Milankovitch cycles. Using a threshold quantile autoregression (TQAR) model, we find strong evidence of nonlinear dynamics. The advantage of the TQAR model is that it allows the dynamics to vary by regime and by quantile. Indeed, dynamics can vary with both past shocks (as captured by regime changes) and with current shocks (as captured by changes across quantiles). As such, the TQAR model can capture the nonlinear dynamics of the climate system without placing *a priori* structural restrictions on the system.

We find statistical evidence that the dynamics of anomalies vary both across regimes and across quantiles. As such, our analysis indicates that linear dynamics would not provide an accurate representation of the evolution of CO_2 concentrations over time. We find that dynamic adjustments are stable when lagged anomalies are in the upper quantile of

the distribution. We also find evidence of local instability, especially when lagged anomalies are around the median and when current anomalies are in the upper quantile. Finally, we find that, overall, dynamic adjustments tend to be slow, indicating that the impact of current shocks can take a long time to develop. This highlights difficulties in using recent information to assess the longer-term impacts of current shocks, making it challenging to evaluate the implications of climate policy.

Using Markov Chain simulations, we investigate the implications of our dynamic analysis for the stationary distribution of anomalies. Due to slow-moving dynamics, we find that the rate of convergence to the stationary distribution to be slow. Our estimates of the shape of the stationary probability function of anomalies shows a departure from Normality, as the distribution is skewed to the left.

The shape of the stationary probability function does not seem to provide information on the presence of tipping points. Note this finding can arise when tipping points are located in specific zones of local instability that depend on past history. For example, our TQAR results showed evidence of local instability when negative shocks for e_{t-1} are followed by positive shocks for e_t. But the different dynamic effects of e_{t-1} and e_t are no longer present in the stationary distribution. Thus, our stationary distribution cannot reveal the local instability uncovered in our TQAR model, which demonstrates the usefulness of our TQAR specification in the investigation of tipping points.

Note that there are scenarios where the presence of tipping points can affect the stationary probabilities. This can occur when the tipping points are in zones of instability involving barriers that are difficult to cross. In this case, falling on one side of an unstable zone would make it difficult to reach the other side, thus creating patterns of dynamic irreversibility. Concerns about possible irreversible dynamics have been raised in the context of climate change. Under irreversibility, one would expect to see the stationary probability function to exhibit low probabilities in zones of local instability (as dynamics would imply escaping from these zones, with a low probability of returning). Our estimated stationary probabilities reported in Fig. 4 do not show such patterns. This seems due to two factors: (1) our estimated zones of instability are not very large (see Fig. 3); and (2) the odds of jumping from one side of an unstable zone to another side appear to be sufficiently high. These two factors reduce the likelihood of observing patterns of irreversibility in the long run, thus contributing to hiding the zones of instability in the stationary probability function. In other words, while our analysis finds evidence of local instability and tipping points, the results reported in Fig. 4 indicate that the associated tipping points do not contribute to generating significant dynamic irreversibility.

Our analysis raises several important questions. How desirable are the different values by CO_2 as climate change affects human welfare? What happens around the tipping points we associated with zones of local instability? What are the circumstances (e.g., policy affecting CO_2 emissions) that may affect the dynamic evolution of CO_2 in these neighborhoods? And what are the possibilities for moving into or out of these neighborhoods? To the extent that one of these neighborhoods could be identified as being "less desirable", trying to avoid it would become an important policy objective. As such, our findings can help inform the current debate about tipping points and climate policy.

Our analysis has uncovered new evidence of nonlinear dynamics and tipping points in the evolution of atmospheric CO_2 concentrations. At this point, it remains unclear what specific aspects of climate dynamics are generating our results. While we have made headway in developing a methodology to examine the nonlinear dynamics of historical climate, there is still much that is left for future research, including looking at reconstructed temperature or atmospheric methane series concomitantly with CO_2 levels; examining behavior at alternative sites; identifying the geophysical processes that can induce nonlinear dynamics in climate evolution; and modeling jointly the paleoclimate in the North and South.

References

Barrett, S., Dannenberg, A., 2012. Climate negotiations under scientific uncertainty. Proc. Natl. Acad. Sci. 109 (43), 17372–17376.
Carpenter, S.R., Kitchell, J.F., Hodgson, J.R., 1985. Cascading trophic interactions and lake productivity. BioScience, 634–639.
Chavas, J.-P., 2015. Modeling population dynamics: a quantile approach. Math. Biosci. 262, 138–146.
Chavas, J.-P., Grainger, C., Hudson, N., 2015. Finding a Fourth "Milankovitch" Cycle in a Reduced-form Regression Framework. Working Paper.
Galvao Jr., A.F., Montes-Rojas, G., Olmo, J., 2011. Threshold quantile autoregressive models. J. Time Ser. Anal. 32 (3), 253–267.
Jansen, E., Overpeck, J., Briffa, K.R., Duplessy, J.C., Joos, F., Masson-Delmotte, V., Ramesh, R., 2007. Paleoclimate. Climate Change 2007: The Physical Science Basis. Working Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change.
Judge, G.G., Hill, R.C., Griffiths, W., Lutkepohl, H., Lee, T.-C., 1988. Introduction to the Theory and Practice of Econometrics.
Kaper, H., Engler, H., 2013. Mathematics and Climate. Siam.
Kelly, D.L., Kolstad, C.D., 1999. Bayesian learning, growth, and pollution. J. Econ. Dyn. Control 23 (4), 491–518.
Koenker, R., 2005. Quantile Regression. Number 38. Cambridge University Press.
Koenker, R., Xiao, Z., 2006. Quantile autoregression. J. Am. Stat. Assoc. 101 (475), 980–990.
Lemoine, D., Traeger, C., 2014. Watch your step: optimal policy in a tipping climate. Am. Econ. J.: Econ. Policy 6 (1), 137–166.
Liu, J., Dietz, T., Carpenter, S.R., Alberti, M., Folke, C., Moran, E., Pell, A.N., Deadman, P., Kratz, T., Lubchenco, J., et al., 2007. Complexity of coupled human and natural systems. Science 317 (5844), 1513–1516.
Lorius, C., Merlivat, L., 1975. Distribution of Mean Surface Stable Isotopes Values in East Antarctica; Observed Changes with Depth in Coastal Area. Technical Report, CEA Centre d'Etudes Nucleaires de Saclay, 91-Gif-sur-Yvette (France). Dept. de Recherche et Analyse.
Lucasius, C.B., Kateman, G., 1993. Understanding and using genetic algorithms. Part 1. Concepts, properties and context. Chemom. Intell. Lab. Syst. 19 (1), 1–33.
Lucasius, C.B., Kateman, G., 1994. Understanding and using genetic algorithms. Part 2. Representation, configuration and hybridization. Chemom. Intell. Lab. Syst. 25 (2), 99–145.
Meyn, S.P., Tweedie, R.L., 2012. Markov Chains and Stochastic Stability. Springer Science & Business Media.
Milanković, M., 1941. Kanon der Erdbestrahlung und seine Anwendung auf das Eiszeitenproblem na.

Nordhaus, W., 2008. A Question of Balance.
Petit, J.-R., Jouzel, J., Raynaud, D., Barkov, N.I., Barnola, J.-M., Basile, I., Bender, M., Chappellaz, J., Davis, M., Delaygue, G., et al., 1999. Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. Nature 399 (6735), 429–436.
Shaphiro, S., Wilk, M., 1965. An analysis of variance test for normality. Biometrika 52, 591–611.
Solomon, S., 2007. Climate Change 2007 – The Physical Science Basis: Working Group I Contribution to the Fourth Assessment Report of the IPCC, vol. 4. Cambridge University Press.
Weitzman, M.L., 2009. On modeling and interpreting the economics of catastrophic climate change. Rev. Econ. Stat. 91 (1), 1–19.
Wright, J.D., Schaller, M.F., 2013. Evidence for a rapid release of carbon at the Paleocene–Eocene thermal maximum. Proc. Natl. Acad. Sci. 110 (40), 15908–15913.